NASCUP: Nucleic Acid Sequence Classification by Universal Probability

Sunyoung Kwon[1], Gyuwan Kim[1], Byunghan Lee[1], Sungroh Yoon[1*], and Young-Han Kim[2*]

[1] *Electrical and Computer Engineering, Seoul National University, Seoul 08826, Korea*
[2] *Electrical and Computer Engineering, University of California, San Diego, La Jolla, CA 92093, USA*

arXiv:1511.04944v1 [q-bio.GN] 16 Nov 2015

Abstract

Motivated by the need for fast and accurate classification of unlabeled nucleotide sequences on a large scale, we propose a new classification method that captures the probabilistic structure of a sequence family as a compact context-tree model and uses it efficiently to test proximity and membership of a query sequence. The proposed *nucleic acid sequence classification by universal probability* (NASCUP) method crucially utilizes the notion of universal probability from information theory in model-building and classification processes, delivering BLAST-like accuracy in orders-of-magnitude reduced runtime for large-scale databases. A comprehensive experimental study involving seven public databases for functional non-coding RNA classification and microbial taxonomy classification demonstrates the advantages of NASCUP over widely-used alternatives in efficiency, accuracy, and scalability across all datasets considered.

1 Introduction

Biological sequence classification plays a key role in various bioinformatics pipelines by revealing the proximity and membership of a sequence to known sequence families [1–5]. Expedited by new sequencing technologies, the amount of biological information is growing rapidly, making the need for having an effective classification method more important than ever. Various state-of-the-art approaches have been proposed, but they often suffer from common limitations in time and space efficiency of training and testing, classification accuracy, and/or scalability.

Existing approaches to sequence classification can be broadly categorized into alignment-based and model/feature-based methods. In alignment-based approaches the similarity between a query sequence and the sequences in a database is measured by sequence-to-sequence alignment, which can be then used to determine the class of the query sequence. Widely used alignment tools include BLAST [6], FASTA [7], BLAT [8], and SOAP [9]. These tools typically exploit the seed-and-expand technique for improving search speed. As the size of the database grows, however, the time demand for database-wide classification often becomes prohibitively costly.

Several types of model/feature-based methods exist. First, probabilistic model-based approaches derive a model from the sequences in a family and perform sequence-to-model comparison in order to classify query sequences. Classification techniques based on hidden Markov models (*e.g.*, HMMER [10, 11] and SAM-T97 [12]), despite their popularity, often suffer from excessive running time caused by iterative training of the model and preprocessing by multiple sequence alignment (MSA). Techniques based on the interpolated context model (ICM) [13] (*e.g.*, GLIMMER [14] and Phymm [15]) show improved time efficiency but often fail to achieve the desired accuracy due to the greedy nature in their search strategy.

Second, k-mer based approaches, such as RDP [16], NBC [17], Kraken [18] and CLARK [19], extract k-mer counts from the data and then perform classification using the acquired k-mer spectra. On the theoretical side, these methods involve somewhat arbitrary metrics for comparison of k-mer spectra, such as Euclidean distances. On the practical side, the need for storing and processing k-mer data often limits the scalability and efficiency of many k-mer based methods.

Compression-based approaches [20,21] constitute the third, and the least popular, type of model/ feature-based methods. Roughly speaking, these approaches utilize the marginal description com-

*To whom correspondence should be addressed: sryoon@snu.ac.kr, yhk@ucsd.edu

plexity of a query sequence when it is compressed as a sequel to database sequences; if the query sequence belongs to the family, then it would be compressed efficiently. The lack of a separate modeling process and the difficulty of compressing all sequences in a family limit the wide adoption of these techniques, not to mention their relatively poor accuracy.

To address the limitations of existing sequence classification methods, we propose the nucleic acid sequence classification by universal probability (NASCUP) method. NASCUP uses a model/feature-based approach with characteristics related to all three types of model/feature-based approaches—based on probabilistic models, k-mers, and compression—but with the following advantages over the existing classification methods:

- **Versus alignment based:** NASCUP does not rely on any sequence-to-sequence comparison, and the classification time of NASCUP is independent of the number of sequences per family, which enables a significant speedup. For gigabase-scale databases, such as Greengenes [22] and SILVA [23], an alignment-based approach would take days or even weeks for database-wide classification, whereas the time demand of NASCUP is orders-of-magnitude lower.

- **Versus k-mer based:** NASCUP utilizes a hierarchical model called *the context-tree model* (CTM), which is constructed from the "contexts" of training sequences using their k-mer information. NASCUP represents the contexts of multiple sequences compactly using a pruned tree structure (not a dense tree used by conventional k-mer based techniques), thus noticeably alleviating the efficiency problem of existing k-mer based approaches.

- **Versus compression based:** For deriving the CTM, NASCUP employs universal probability assignments as in the method of context-tree maximizing compression [24, 25]. This allows for efficient and effective modeling and classification of multiple sequences in two stages, improving significantly upon the performance and scalability of compression based approaches even though NASCUP shares much of their philosophical foundations.

- **Versus probabilistic model based:** NASCUP runs faster and returns more accurate classification results than HMM-based and ICM-based alternatives. Furthermore, NASCUP does not need any MSA preprocessing, which allows additional performance boosts over the HMM approaches that normally require MSA preprocessing. Compared with ICM, which also utilizes context trees, NASCUP implicitly enumerates all instances of a context tree of a given depth and traverses them so that the most probable CTM can be found. This principled approach of tree building in NASCUP pays off with its much improved accuracy. Even so, the running time of NASCUP typically remains lower than that of ICM.

We have performed a set of comprehensive experiments that test NASCUP with a variety of real sequence datasets for functional non-coding RNA classification and microbial taxonomy classification. Our experimental results confirm the effectiveness of NASCUP and its advantages over existing state-of-the-art in terms of time and space efficiency and classification accuracy.

2 Proposed Method

NASCUP aims to capture the probabilistic structure of each nucleotide sequence family by relative frequencies of the symbols `A`,`C`,`G`,`T`(`U`) in the constituent sequences. The most straightforward approach to distilling this information is to count k-mers [16–19], or equivalently, counting the symbols conditioned on depth $d = k - 1$ Markov states or *contexts*. This naïve approach, however, has two serious flaws. First, there are 4^d potential contexts (which is a huge number for a reasonable depth d), many of which may not appear at all. Second, even those that appear a few times may not contribute to any statistically meaningful structure and instead cause unnecessary overfitting.

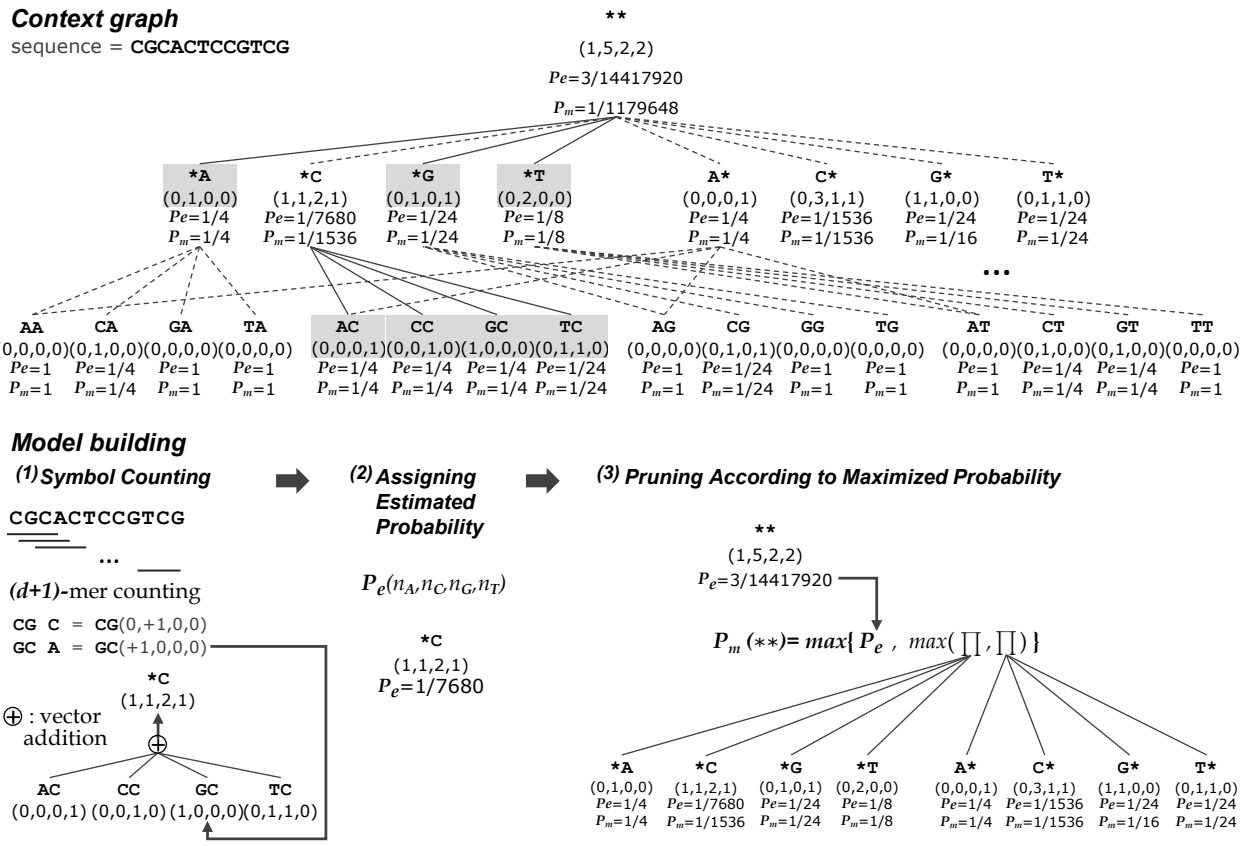

Fig. 1: Context graph with depth 2 and the model-building process for a single database sequence X=CGCACTCCGTCG [solid and dotted lines: selected and pruned contexts, respectively; the set of shaded regions: the optimal CTM found].

Similar to ICM [13] and its variations [14, 15], NASCUP overcomes these flaws by employing variable-order Markov models or more general context-tree models. Such context tree models tend to capture the probabilistic structure of sequences more accurately with far fewer contexts. As a simple analogy, in the English language, symbol 'Q' is most likely followed by 'U' and what precedes 'Q' does not affect this behavior in any significant manner. Thus, NASCUP starts with a context graph (as shown in Fig. 1 for $d = 2$) that consists of nodes representing all potential contexts $\{\mathtt{A}, \mathtt{C}, \mathtt{G}, \mathtt{T}, *\}^d$, where $*$ denotes "don't care," and builds a context tree by keeping only statistically meaningful branches of the context graph.

The main feature that distinguishes NASCUP with other model-based methods utilizing context trees is that NASCUP chooses the context tree (among all possible context trees that may arise from the context graph) that has the maximum *universal* probability. This model-building method closely resembles the context-tree maximizing data compression algorithm [24] (see also [25]), which is known to achieve the optimal compression performance [26], as well as tree-based decision algorithms [27,28]. NASCUP, however, does not account for the description complexity of the tree model itself, which is crucial in compression or tree-based decision making, but is irrelevant in classification.

Given a new sequence whose family membership is unknown, NASCUP compares the (conditional) probabilities of the sequence given context trees for the sequence families. These probabilities are once again computed according to universal probability assignments, and the family with highest probability is selected. In the language of data compression, this is equivalent to choosing the best source code among many codes attuned to different probabilities. In this sense, NASCUP can be viewed as a refinement of compression-based classification methods [20, 21] that uses the optimal source code using universal probability.

In the following, we describe how NASCUP works in two stages—model building and classification.

2.1 Model Building

NASCUP builds a context-tree model (CTM) for each sequence family in the given dataset.

2.1.1 Step 1: Symbol Counting for Each Context

NASCUP initially counts the nucleotide symbols A,C,G,T that follow each of length-d contexts in $\{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}^d$ by enumerating all k-mers ($k = d+1$) from all sequences in the family. These counts are then merged bottom–up along the context graph of depth d as shown in Figure 1.

The context graph has as its nodes 5^d contexts that include the "don't care" ($*$) symbol in addition to $\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}$, with $**\cdots*$ being the root (depth 0) and physical contexts from $\{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}^d$ being 4^d leaves at the bottom (depth d). Each context from depth 0 to $d-1$ is branched into one or more sets of children in the next depth by replacing each $*$ symbol with a set of $\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}$. The root has d sets of children. More generally, a context with j $*$ symbols has j sets of children. For example, context $*\mathtt{A}$ at depth 1 is branched into one children set $\mathtt{AA},\mathtt{CA},\mathtt{GA},\mathtt{TA}$ at depth 2, and $\mathtt{AA}$ at depth 2 is, in turn, connected to two parents $*\mathtt{A}$ and $\mathtt{A}*$. The counts of the symbols at each context $s \in \{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T},*\}^d$ is denoted as $n_i(s)$, $i \in \{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}$, which are collectively written as a 4-dim count vector $\mathbf{n}(s)$, *e.g.*,

$$\mathbf{n}(**) = \mathbf{n}(*\mathtt{A}) + \mathbf{n}(*\mathtt{C}) + \mathbf{n}(*\mathtt{G}) + \mathbf{n}(*\mathtt{T}) = \mathbf{n}(\mathtt{A}*) + \mathbf{n}(\mathtt{C}*) + \mathbf{n}(\mathtt{G}*) + \mathbf{n}(\mathtt{T}*).$$

In case the best variable-order Markov model (VMM) is needed in place of the best CTM, NASCUP restricts the initial context graph to a subgraph (tree) on Markov contexts, namely, the contexts in $\{*\}^j \times \{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}^{d-j}$ for some j.

2.1.2 Step 2: Assigning Estimated Probabilities to Each Context

NASCUP assigns zeroth-order (or elementary) probabilities to each context s based on its count vector $\mathbf{n}(s)$. These probability assignments follow well-established alternatives in the data compression and machine learning literature. In the following, we omit s whenever it is clear from the context.

- Krichevski–Trofimov (KT) [29]: This probability estimate is a Dirichlet mixture of independent and identically distributed (i.i.d.) random processes over the quaternary alphabet, which is known to achieve the minimum compression redundancy for i.i.d. processes with unknown parameters and is used widely in universal data compression [30]:

$$P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}}) = \frac{\prod_{i\in\{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}} \prod_{j_i=1}^{n_i} (j_i - 1/2)}{\prod_{j=1}^{n_{\mathtt{A}}+n_{\mathtt{C}}+n_{\mathtt{G}}+n_{\mathtt{T}}} (j+1)}, \tag{1}$$

  where an empty product is set to 1 by convention.

- Zero redundancy (ZR) [31,32]: This variant of the KT estimate, denoted by $P_e^{\mathrm{ZR}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}})$, assigns more weights to the cases with zero counts (more emphasis on structural properties of the process):

$$\begin{aligned}
\frac{1}{16}\big[&2P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}}) + P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}})\delta(n_{\mathtt{T}}) + P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{T}})\delta(n_{\mathtt{G}}) \\
&+ P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{G}}, n_{\mathtt{T}})\delta(n_{\mathtt{C}}) + P_e^{\mathrm{KT}}(n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}})\delta(n_{\mathtt{A}}) \\
&+ P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}})\delta(n_{\mathtt{G}}, n_{\mathtt{T}}) + P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{G}})\delta(n_{\mathtt{C}}, n_{\mathtt{T}}) + P_e^{\mathrm{KT}}(n_{\mathtt{A}}, n_{\mathtt{T}})\delta(n_{\mathtt{C}}, n_{\mathtt{G}}) \\
&+ P_e^{\mathrm{KT}}(n_{\mathtt{C}}, n_{\mathtt{G}})\delta(n_{\mathtt{A}}, n_{\mathtt{T}}) + P_e^{\mathrm{KT}}(n_{\mathtt{C}}, n_{\mathtt{T}})\delta(n_{\mathtt{A}}, n_{\mathtt{G}}) + P_e^{\mathrm{KT}}(n_{\mathtt{G}}, n_{\mathtt{T}})\delta(n_{\mathtt{A}}, n_{\mathtt{C}}) \\
&+ \delta(n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}}) + \delta(n_{\mathtt{A}}, n_{\mathtt{G}}, n_{\mathtt{T}}) + \delta(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{T}}) + \delta(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}})\big],
\end{aligned} \tag{2}$$

  where $\delta(\cdot) = 1$ if its argument is a zero vector; otherwise $\delta(\cdot) = 0$. Here $P_e^{\mathtt{KT}}(\cdot)$ is defined

similarly as in (1) when taking fewer than four arguments. For example,

$$P_e^{\mathtt{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}) = \frac{\prod_{j_{\mathtt{A}}=1}^{n_{\mathtt{A}}}(j_{\mathtt{A}} - 1/2)\prod_{j_{\mathtt{C}}=1}^{n_{\mathtt{C}}}(j_{\mathtt{C}} - 1/2)}{\prod_{j=1}^{n_{\mathtt{A}}+n_{\mathtt{C}}} j}, \tag{3}$$

$$P_e^{\mathtt{KT}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}}) = \frac{\prod_{j_{\mathtt{A}}=1}^{n_{\mathtt{A}}}(j_{\mathtt{A}} - 1/2)\prod_{j_{\mathtt{C}}=1}^{n_{\mathtt{C}}}(j_{\mathtt{C}} - 1/2)\prod_{j_{\mathtt{G}}=1}^{n_{\mathtt{G}}}(j_{\mathtt{G}} - 1/2)}{\prod_{j=1}^{n_{\mathtt{A}}+n_{\mathtt{C}}+n_{\mathtt{G}}}(j + 1/2)}. \tag{4}$$

- Maximum likelihood (ML): The maximum likelihood probability estimate is used in many applications and we include it as a reference method:

$$P_e^{\mathrm{ML}}(n_{\mathtt{A}}, n_{\mathtt{C}}, n_{\mathtt{G}}, n_{\mathtt{T}}) = \prod_{i\in\{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}} \left(\frac{n_i}{n_{\mathtt{A}} + n_{\mathtt{C}} + n_{\mathtt{G}} + n_{\mathtt{T}}}\right)^{n_i}, \tag{5}$$

  although it does not technically induce a valid probability distribution over sequences of a given length (the sum of probabilities is greater than 1) and it often leads to overfitting (especially when some symbol never appears). As a simple remedy, one can add a pseudocount [33], say, $c = 1/2$, to the count vector $\mathbf{n}$. Note that this pseudocount adjustment can be viewed also as an approximation of the KT estimate using a homogenous (instead of running) product.

2.1.3 Step 3: Pruning Contexts According to Maximized Probabilities

Starting with the root node in the context graph in a depth-first manner, NASCUP computes the maximized probability $P_m(s)$ for each context s as

$$P_m(s) = \max\Big\{P_e(s), \max_{\text{children sets of } s} \prod_{\text{child } s' \text{ in each set}} P_m(s')\Big\}, \tag{6}$$

where $P_e(s) = P_e(n_{\mathtt{A}}(s), n_{\mathtt{C}}(s), n_{\mathtt{G}}(s), n_{\mathtt{T}}(s))$ denotes the estimated probability of the context s (according to one of the four estimated probability assignment methods in step 2) and the inner maximum is over different children sets of s. If s is a leaf (at depth d), then $P_m(s) = P_e(s)$. As an example, for $d = 2$,

$$P_m(\mathtt{**}) = \max\{P_e(\mathtt{**}),\ P_m(\mathtt{*A})P_m(\mathtt{*C})P_m(\mathtt{*G})P_m(\mathtt{*T}),\ P_m(\mathtt{A*})P_m(\mathtt{C*})P_m(\mathtt{G*})P_m(\mathtt{T*})\},$$
$$P_m(\mathtt{*A}) = \max\{P_e(\mathtt{*A}),\ P_e(\mathtt{AA})P_e(\mathtt{CA})P_e(\mathtt{GA})P_e(\mathtt{TA})\}.$$

Each children set that does not attain the maximum is pruned. In case of a tie between a context and a children set, we choose the context for the sparsity of the resulting model. Upon completion of this maximizing and pruning step, the remaining nodes of the graph is either a leaf or a non-leaf with a single non-overlapping set of four children. This tree and the count vectors of its nodes (contexts) form the context-tree model (CTM) for the given sequence family. We remark that our probability maximization method is different from the context-tree maximizing method in data compression [24, 25] in that we do not include the description complexity of subtrees.

2.2 Classification

Assume that there are K known families of sequences in the database and that NASCUP has already constructed context-tree models for them. Given an unlabeled input sequence $\mathbf{x}$, NASCUP computes the membership scores of the sequence for the K context-tree models in three steps.

2.2.1 Step 1. Symbol Counting for Each Context

As in the modeling stage, NASCUP performs $(d+1)$-mer enumeration on $\mathbf{x}$ to generate count vectors $\mathbf{m}(s) = (m_{\mathtt{A}}(s), m_{\mathtt{C}}(s), m_{\mathtt{G}}(s), m_{\mathtt{T}}(s))$ for each possible context s of length d.

2.2.2 Step 2. Probability Computation for Each Model

Fix a context-tree model M characterized by a context tree (a set S of contexts) and the count vectors $\mathbf{n}(s)$ for $s \in S$. NASCUP computes the conditional probability of $\mathbf{x}$ given M as

$$P(\mathbf{x}|M) = \prod_{s \in S} P_c(\mathbf{m}(s)|\mathbf{n}(s)), \tag{7}$$

where the classification probability $P_c(\mathbf{m}(s)|\mathbf{n}(s))$ is computed according to one of the following:

- KT: The classification probability is computed as a conditional probability (or equivalently, the Bayesian update) of the new sequence given the training sequences under the KT probability assignment in (1):
$$P_c^{\mathrm{KT}}(\mathbf{m}|\mathbf{n}) = \frac{P_e^{\mathrm{KT}}(\mathbf{m}+\mathbf{n})}{P_e^{\mathrm{KT}}(\mathbf{n})}. \tag{8}$$
- ZR: The classification probability is computed as in (2), except that $P_c^{\mathrm{KT}}(\cdot|\cdot)$ in (8) is used instead of $P_e^{\mathrm{KT}}(\cdot)$ in (1), (3), and (4).
- ML: The classification probability uses the model count vector with pseudocount adjustment ($c = 1/2$):
$$P_c^{\mathrm{ML}}(\mathbf{m}|\mathbf{n}) = \prod_{i \in \{\mathtt{A},\mathtt{C},\mathtt{G},\mathtt{T}\}} \left(\frac{n_i + c}{n_{\mathtt{A}} + n_{\mathtt{C}} + n_{\mathtt{G}} + n_{\mathtt{T}} + 4c} \right)^{m_i}. \tag{9}$$

2.2.3 Step 3. Finding the Best Model

NASCUP computes the statistical validity of the assignment of $\mathbf{x}$ to each model M by computing the log-odds score

$$\log_4 \frac{P(\mathbf{x}|M)}{(1/4)^{l-d}}, \tag{10}$$

where l denotes the length of $\mathbf{x}$. This score compares $P(\mathbf{x}|M)$ and the probability of the i.i.d. uniform sequence (null model). This can be equivalently viewed as the compression efficiency of $\mathbf{x}$ under a universal source code trained with the sequences in the model; when normalized by the sequence length, it captures the compression ratio (with 0% being no compression and 100% being complete compression). NASCUP assigns the sequence $\mathbf{x}$ to the family whose model M has the highest log-odds score, provided that its score is above some threshold.

3 Experimental Results

3.1 Setup and Datasets

We implemented NASCUP in C++ and measured its performance in terms of runtime (*i.e.*, modeling time plus classification time), and classification accuracy with 10-fold cross validation, using a linux machine (Ubuntu 12.04, 2.2 GHz Intel Xeon E5-4620, and 512 GB memory) without any parallelization. For comparison, we also prepared the implementations of BLAST (blastall 2.2.25 in C++), HMMER (3.1.b2 in C), ICM (PhymmBL v4.0 in C++), and gzip (v1.4 in C).

We used seven datasets of two types—functional non-coding RNAs and microbial taxonomy (Table 1). For the former, we used Rfam [34], a database of functional RNA families of rRNAs, tRNAs, miRNAs, siRNAs, and cis-regulatory elements. For the latter, we used databases of rRNA sequences (RDP [35], Greengenes [22], and SILVA [23]) and pyrosequencing datasets (Artificial [36] and Divergent [36]). SILVA consisted of two subsets, one from large subunit (LSU) 23S and 28S rRNA

Table 1: Details of the datasets used in the experiments

classification category of the dataset		ID	dataset name	release (version)	# families*	# total† sequences	sequence length	ground truth	source
functional non-coding RNA		RF	Rfam	11.0	1,320	170,881	20–1,875	accession	[34]
microbial taxonomy	rRNA database (16S, 18S, 23S/28S)	RD	RDP	10.0	134	3,838	320–1,833		[35,36]
		GG	Greengenes	13.5	464	23,142	1,254–2,146	taxonomy	[22]
		SS	SILVA-SSU	119.1	313	17,625	902–3,749	(genus level)	[23]
		SL	SILVA-LSU	119	107	4,593	1,900–4,954		[23]
	pyrosequencing data (16S rRNA)	AR	Artificial	-	60	44,407	40–294	reference	[37]
		DV	Divergent	-	23	55,466	38–521	sequences	[37]

* the number of families with more than 10 sequences † the total number of sequences after the preprocessing

sequences and the other from small subunit (SSU) 16S and 18S sequences. The two pyrosequencing datasets contained mixtures of microbial 16S rRNA sequenced with the 454 technology.

As preprocessing, we ignored the sequences containing symbols other than `ACGT(U)` and, for 10-fold cross validation, excluded the sequence families with less than ten sequences. Since some methods compared could not handle large data in reasonable time, we reduced the data as follows: First, we limited the maximum number of sequences in each family to 2,000 by random sampling. For Greengenes and SILVA, we further modified the sampled sequences by cd-hit-est [37] using a 97% of similarity [38] threshold. For RDP, we downloaded a public RDP dataset modified by other researchers using mothur [39] and preprocessed it as above.

The source code and test data will be available at http://data.snu.ac.kr/nascup.

3.2 Performance Evaluation: Nuts and Bolts of NASCUP

We first evaluated NASCUP by altering its components and depth d to find their optimal combination. We tested a total of 216 combinations: nine variations in modeling, four variations in classification, and six values of $d = 3, 4, \ldots, 8$.

We used the KT, ZR, and ML[‡] estimators implemented as explained in Section 2. As an additional option for model building, we implemented KMER that performs simple k-mer counting and uses the result as the model, as is usually the case with conventional k-mer methods. As an option for classification, we also implemented EUC, which classifies a query sequence based on composite Euclidean distances between context count vectors of the query and database sequences.

3.2.1 Comparison of Different Model-Building Options in Time and Space

For each of the nine methods used in model building, Fig. 2(a) shows the runtime normalized by that of KMER (depth 3). For the same depth, the variation of runtime across methods was negligible, except for ZR, which involves running multiple rounds of the KT estimation. More importantly, there was no noticeable difference in runtime to build a CTM or a VMM, even though a VMM is only one of the $d!$ possible structures considered for building a CTM. This result suggests the effectiveness of pruning explained in Section 2.1.3, which successfully avoided enumerating most of the $O(d!4^{d+1})$ paths NASCUP would traverse in the worst case. On the other hand, the modeling time grew rapidly as we increased d, since the size of the context graph increased exponentially in d, directly affecting the time demand for k-mer counting and computation of P_e and P_m.

Fig. 2(b) shows, for each configuration, the number of contexts in the model (*i.e.*, the number of leaf nodes). This quantity affects the overall performance in two aspects: Having fewer contexts will be helpful for reducing the memory usage and for expediting classification. For $d < 6$, there was no noticeable difference among the methods tried. From $d = 6$ and on, however, we observed significant reduction in the number of contexts by using KT/ZR/ML over KMER.

[‡]ML and ML(1/2) represents the ML estimator without and with pseudocount $c = 1/2$, respectively.

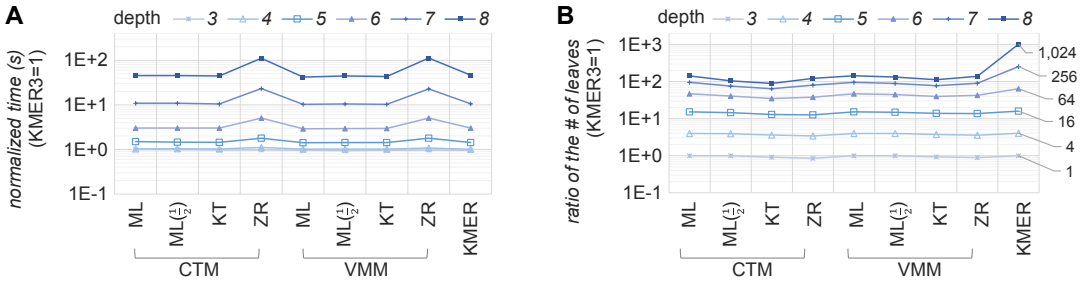

Fig. 2: Comparison of different options for modeling: (A) Runtime of different methods normalized by that of KMER (depth 3). (B) The number of contexts in the model normalized by that of KMER (depth 3). Each data point represents the geometric mean of the runtime or # contexts measured over the seven datasets listed in Table 1.

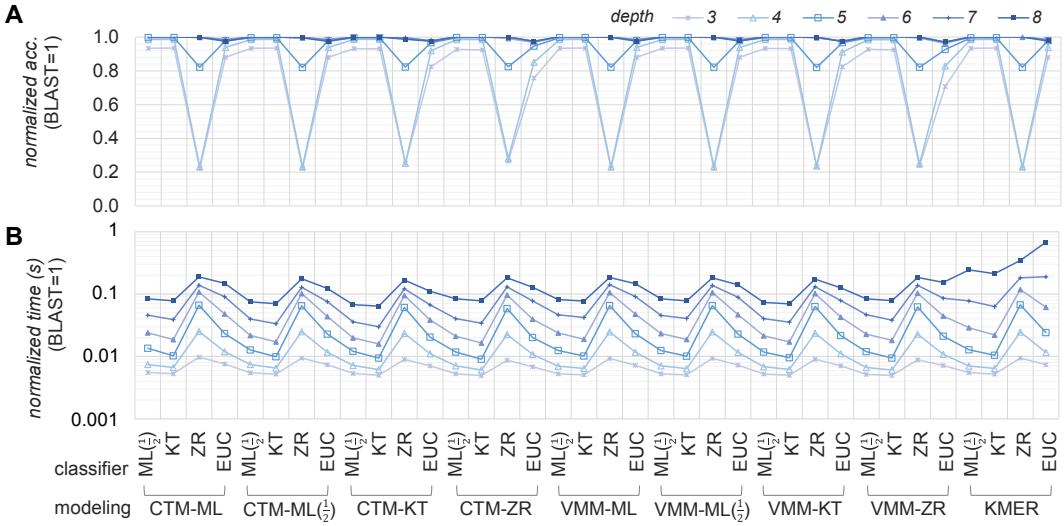

Fig. 3: Comparison of different combinations of modeling and classification options: (A) Accuracy. (B) Classification time. Each data point represents the geometric mean of the accuracy or runtime measured over the seven datasets listed in Table 1.

3.2.2 Comparison of Different Classification Options in Accuracy and Time

For different combinations of modeling and classification options, Fig. 3(a) shows the classification accuracy of each combination normalized by that of BLAST-based nearest-neighbor classification (abbreviated as BLAST). Most notably, as long as $d \geq 5$, the normalized accuracy of all combinations remained near one without being affected by increasing d, except for the combinations using ZR for classification. This indicates that we can attain satisfactory performance without increasing d too much (which would result in time/space inefficiency and in potential overfitting).

The impact of d on time efficiency is evident as shown in Fig. 3(b), which presents the classification time of different combinations normalized by that of BLAST. Regardless of the combination used, the time grew rapidly as d increased. Nonetheless, most combination of the NASCUP components, with ZR combinations as notable exceptions, ran significantly faster than BLAST.

Examining the accuracy and classification time results together suggests that using $d = 6$ or $d = 7$ would best balance accuracy and classification time.

3.3 Performance Comparison of NASCUP with Other Classification Methods

Figs. 4(a) and (b) respectively shows the total runtime (*i.e.*, modeling plus classification) and the classification time of NASCUP (labeled VMM-KT/KT6) and the other methods. Since modeling

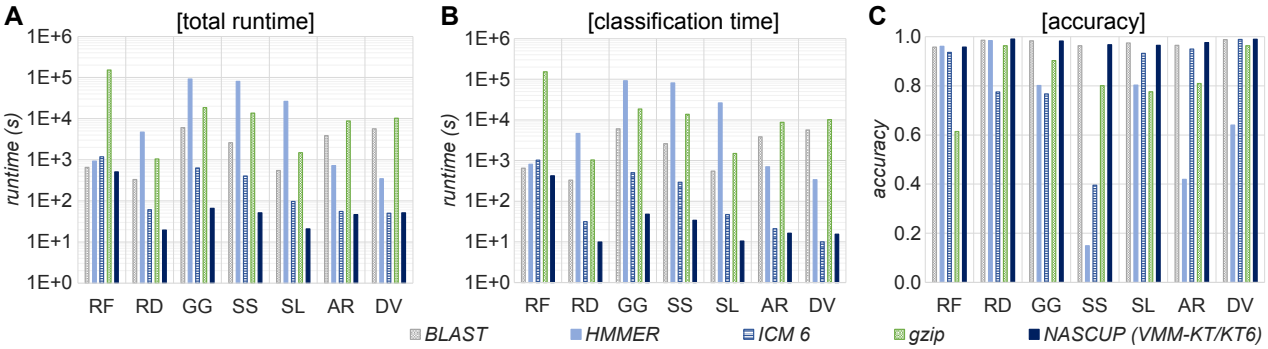

Fig. 4: Comparison of NASCUP with alternatives [ICM6: ICM with depth 6; VMM-KT/KT6: VMM-KT for modeling and KT for classification (both with depth 6)]: (a) Total runtime (modeling + classification). (b) Classification time only. (c) Accuracy of different methods for the seven datasets listed in Table 1.

Table 2: Performance of the best tool for each dataset and relative performance of others

	data	RF	RD	GG	SS	SL	AR	DV
best[‡]	method	KMER/ZR7	CTM-ML(c)/ZR7, CTM-ZR/ZR 6	BLAST	CTM-ZR/ML(c)7	KMER/ZR8	KMER/EUC7	KMER/ZR8
	accuracy	**0.9726**	**0.9911**	**0.9835**	**0.9689**	**0.9800**	**0.9786**	**0.9910**
perf. ratio[¶]	BLAST	98.51%	99.42%	100.00%	99.41%	99.42%	98.58%	99.72%
	HMMER	98.79%	99.26%	81.46%	15.41%	81.96%	42.83%	64.60%
	ICM 6	96.24%	78.23%	78.03%	40.79%	95.09%	97.08%	99.76%
	gzip	64.48%	97.19%	91.81%	82.67%	79.20%	82.69%	97.20%
	VMM-KT/KT6	98.45%	99.89%	99.87%	99.88%	98.47%	99.76%	99.85%

[‡] the best method and its accuracy for each dataset

[¶] performance ratio (in %) of each tool to the best.

occurs once but classification runs multiple times, we plotted the classification time separately. In our experiments, the total time was indeed dominated by classification because of the large size of the data used. We also confirmed that NASCUP ran significantly (often by orders of magnitude) faster for most of the data. The second best was ICM (depth 6), another CTM-based approach. The performance of ICM, however, did not match that of NASCUP in terms of accuracy, as demonstrated in Fig. 4(c) that shows the accuracy of NASCUP and the others. NASCUP and BLAST maintained accuracy higher than 0.95 for all datasets, whereas the accuracy of HMMER and ICM varied depending on the data used. ICM showed accuracy lower than 0.8 for three datasets. To compare the accuracy more closely, we use Table 2. The upper part shows the best tool for each dataset and its accuracy, while the lower part shows the performance ratio (in %) of each tool to the best. NASCUP ranked top for six out of the seven datasets, while remaining competitive for the other dataset as well (99.87% performance of the best tool).

Given that NASCUP and BLAST achieved similar accuracy, we further compared the two in terms of scalability. We first tested how the modeling and classification time of NASCUP was affected as the database size grew. To this end, we gradually increased the number of sequences sampled from the full GG data (from 117,837 to 836,043 sequences) and measured the modeling time and classification time (normalized per sequence and per family), as shown in Fig. 5(a). As expected, the modeling time grew as the size of dataset increases. On the other hand, the classification time did not grow as fast as the modeling time. This is because the classification time of NASCUP is independent of the number of sequences in a family, once the modeling has been completed. This feature makes NASCUP scalable to large data, as shown in Fig. 5(b), which shows that the total runtime of NASCUP is orders of magnitude lower than that of BLAST regardless of the data size.

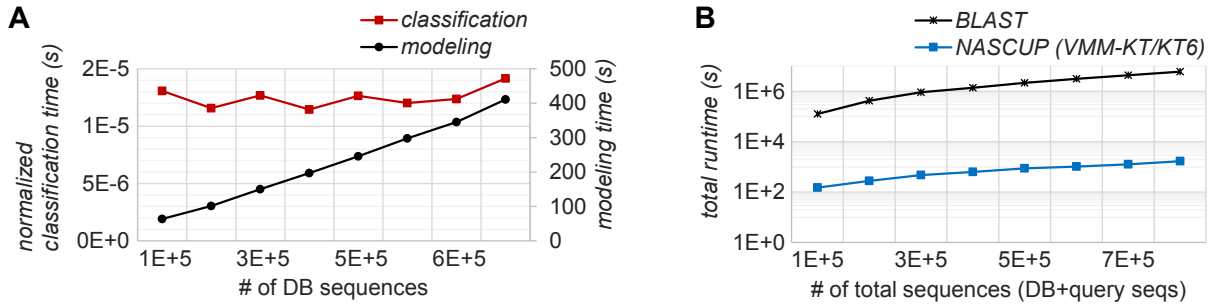

Fig. 5: Scalability tests using the Greengenes [22] database and its subsets: (a) Modeling time vs normalized classification time of NASCUP. (b) Variation of total runtime of NASCUP and BLAST over the data size.

The bottom line from our experiments can be summarized as follows: NASCUP and BLAST were the two *universal* tools that showed satisfactory performance on all the datasets. However, NASCUP was the only *scalable* tool that could process large-scale data efficiently and accurately.

4 Discussion

The proposed NASCUP method performs model-based classification using the universal probability assignment for context-tree models from information theory. Thus, it combines the merits of existing model-based methods and compression-based methods, which is confirmed by its highly competitive performance in diverse datasets in terms of time efficiency and classification accuracy.

As a model-based method, NASCUP can be easily adapted to test similarity between families or a membership of a sequence or a sequence family in a group of families. We expect that this modified version of NASCUP can be applied to problems with inherent hierarchical structures such as taxonomy classification and phylogenetic tree construction. A phylogenetic tree is typically constructed by measuring sequence similarities using heuristics such as progressive alignment and then by applying linkage methods to produce a dendrogram structure. In this conventional approach, results may vary (often dramatically) according to which linkage method is used. In contrast, the modified version of NASCUP would allow direct model-to-model comparison, alleviating the aforementioned limitations of alignment–linkage methods. Going one step further, we expect that the modified NASCUP, fully utilizing the simple Bayesian update structure of universal probability assignments in the classification stage, will be applicable to incremental clustering, in which an initial model is constructed from close sequences and then is incrementally augmented by adding additional sequences.

More technical aspects of NASCUP can be improved in the following directions, providing the end users with more flexibility. First, instead of the current classification scoring scheme of log-odds against random data, which is effective in many cases but cannot present an absolute criterion, we can adopt a more standardized scoring scheme. One such scheme can be found in operational taxonomic unit (OTU) [40] binning, whereby a widely used threshold is 97% similarity [38] at the species level. Second, instead of unidirectional contexts (that precede symbols), we can use bidirectional contexts (that precede and succeed symbols). More generally, the directionality of nucleotide sequences or the lack thereof can be incorporated into modeling and classification stages. Third, the maximum depth of context-tree models can be adjusted more flexibly in a data-dependent manner. A similar technique has been developed in the context of data compression [41], which is expected to be applicable in our problem. The use of bidirectional contexts and adaptive maximal depths is expected to boost accuracy further at some cost of time efficiency.

References

[1] M. H. van Regenmortel, C. M. Fauquet, D. H. Bishop, E. Carstens, M. Estes, S. Lemon, J. Maniloff, M. Mayo, D. McGeoch, C. Pringle, *et al.*, *Virus taxonomy: classification and nomenclature of viruses. Seventh report of the International Committee on Taxonomy of Viruses.* Academic Press, 2000.

[2] J. Lu, G. Getz, E. A. Miska, E. Alvarez-Saavedra, J. Lamb, D. Peck, A. Sweet-Cordero, B. L. Ebert, R. H. Mak, A. A. Ferrando, *et al.*, "MicroRNA expression profiles classify human cancers," *nature*, vol. 435, no. 7043, pp. 834–838, 2005.

[3] S. Griffiths-Jones, S. Moxon, M. Marshall, A. Khanna, S. R. Eddy, and A. Bateman, "Rfam: annotating non-coding RNAs in complete genomes," *Nucleic acids research*, vol. 33, no. suppl 1, pp. D121–D124, 2005.

[4] P. Yap, J. Kolberg, and M. Urdea, "Classification of hepatitis C virus into six major genotypes and a series of subtypes by phylogenetic analysis of the ns-5 region," *Journal of general virology*, vol. 74, pp. 2391–2399, 1993.

[5] R. L. Tatusov, D. A. Natale, I. V. Garkavtsev, T. A. Tatusova, U. T. Shankavaram, B. S. Rao, B. Kiryutin, M. Y. Galperin, N. D. Fedorova, and E. V. Koonin, "The COG database: new developments in phylogenetic classification of proteins from complete genomes," *Nucleic acids research*, vol. 29, no. 1, pp. 22–28, 2001.

[6] S. F. Altschul, W. Gish, W. Miller, E. W. Myers, and D. J. Lipman, "Basic local alignment search tool," *Journal of molecular biology*, vol. 215, no. 3, pp. 403–410, 1990.

[7] W. R. Pearson, "[5] rapid and sensitive sequence comparison with FASTP and FASTA," *Methods in enzymology*, vol. 183, pp. 63–98, 1990.

[8] W. J. Kent, "BLAT—the BLAST-like alignment tool," *Genome research*, vol. 12, no. 4, pp. 656–664, 2002.

[9] R. Li, Y. Li, K. Kristiansen, and J. Wang, "SOAP: short oligonucleotide alignment program," *Bioinformatics*, vol. 24, no. 5, pp. 713–714, 2008.

[10] T. J. Wheeler and S. R. Eddy, "nhmmer: DNA homology search with profile HMMs," *Bioinformatics*, p. btt403, 2013.

[11] S. R. Eddy, "Accelerated profile HMM searches," *PLoS Comput Biol*, vol. 7, no. 10, p. e1002195, 2011.

[12] K. Karplus, C. Barrett, and R. Hughey, "Hidden Markov models for detecting remote protein homologies.," *Bioinformatics*, vol. 14, no. 10, pp. 846–856, 1998.

[13] S. L. Salzberg, A. L. Delcher, S. Kasif, and O. White, "Microbial gene identification using interpolated Markov models," *Nucleic Acids Research*, vol. 26, no. 2, pp. 544–548, 1998.

[14] A. L. Delcher, D. Harmon, S. Kasif, O. White, and S. L. Salzberg, "Improved microbial gene identification with GLIMMER," *Nucleic Acids Research*, vol. 27, no. 23, pp. 4636–4641, 1999.

[15] A. Brady and S. L. Salzberg, "Phymm and PhymmBL: Metagenomic phylogenetic classification with interpolated Markov models," *Nature Methods*, vol. 6, no. 9, pp. 673–676, 2009.

[16] J. R. Cole, Q. Wang, E. Cardenas, J. Fish, B. Chai, R. J. Farris, A. Kulam-Syed-Mohideen, D. M. McGarrell, T. Marsh, G. M. Garrity, *et al.*, "The Ribosomal Database Project: improved alignments and new tools for rRNA analysis," *Nucleic acids research*, vol. 37, no. suppl 1, pp. D141–D145, 2009.

[17] G. Rosen, E. Garbarine, D. Caseiro, R. Polikar, and B. Sokhansanj, "Metagenome Fragment Classification using N-Mer Frequency Profiles," *Advances in bioinformatics*, vol. 2008, 2008.

[18] D. E. Wood and S. L. Salzberg, "Kraken: ultrafast metagenomic sequence classification using exact alignments," *Genome Biol*, vol. 15, no. 3, p. R46, 2014.

[19] R. Ounit, S. Wanamaker, T. J. Close, and S. Lonardi, "CLARK: fast and accurate classification of metagenomic and genomic sequences using discriminative k-mers," *BMC genomics*, vol. 16, no. 1, p. 236, 2015.

[20] A. Kocsor, A. Kertész-Farkas, L. Kaján, and S. Pongor, "Application of compression-based distance measures to protein sequence classification: a methodological study," *Bioinformatics*, vol. 22, no. 4, pp. 407–412, 2006.

[21] D. Loewenstern, H. Hirsh, P. Yianilos, and M. Noordewier, *DNA sequence classification using compression-based induction*. Citeseer, 1995.

[22] T. Z. DeSantis, P. Hugenholtz, N. Larsen, M. Rojas, E. L. Brodie, K. Keller, T. Huber, D. Dalevi, P. Hu, and G. L. Andersen, "Greengenes, a chimera-checked 16S rRNA gene database and workbench compatible with ARB," *Applied and environmental microbiology*, vol. 72, no. 7, pp. 5069–5072, 2006.

[23] C. Quast, E. Pruesse, P. Yilmaz, J. Gerken, T. Schweer, P. Yarza, J. Peplies, and F. O. Glöckner, "The SILVA ribosomal RNA gene database project: improved data processing and web-based tools," *Nucleic acids research*, p. gks1219, 2012.

[24] F. M. J. Willems, Y. M. Shtarkov, and T. J. Tjalkens, "Context weighting: General finite context sources," in *14th Symp. Inf. Theory Benelux*, (Veldhoven, The Netherlands), pp. 120–127, May 1993.

[25] R. Nohre, *Some topics in descriptive complexity*. Ph.D. thesis, Linköping University, Linköping, Sweden, 1994.

[26] J. Rissanen, "Universal coding, information, prediction, and estimation," *Information Theory, IEEE Transactions on*, vol. 30, no. 4, pp. 629–636, 1984.

[27] J. R. Quinlan and R. L. Rivest, "Inferring decision trees using the minimum description lenght principle," *Information and computation*, vol. 80, no. 3, pp. 227–248, 1989.

[28] M. Mehta, J. Rissanen, and R. Agrawal, "Mdl-based decision tree pruning," in *KDD*, pp. 216–221, Aug. 1995.

[29] R. E. Krichevsky and V. K. Trofimov, "The performance of universal encoding," *Information Theory, IEEE Transactions on*, vol. 27, no. 2, pp. 199–207, 1981.

[30] T. M. Cover and J. A. Thomas, *Elements of Information Theory*. New York: Wiley, second ed., 2006.

[31] T. J. Tjalkens, P. A. J. Volf, and F. M. J. Willems, "A context-tree weighting method for text generating sources," in *Proc. Data Compression Conf.*, (Snowbird, UT), p. 472, IEEE, Mar. 1997.

[32] M. Rashid and T. Kawabata, "Theoretical analysis of a zero-redundancy estimator with a finite window for memoryless source," in *Proc. IEEE Inf. Theory Works.*, (Rotorua, New Zealand), pp. 171–175, Aug. 2005.

[33] S. Vucetic and Z. Obradovic, "Classification on data with biased class distribution," in *Machine Learning: ECML 2001*, pp. 527–538, Springer, 2001.

[34] S. Griffiths-Jones, A. Bateman, M. Marshall, A. Khanna, and S. R. Eddy, "Rfam: an RNA family database," *Nucleic acids research*, vol. 31, no. 1, pp. 439–441, 2003.

[35] J. R. Cole, Q. Wang, J. A. Fish, B. Chai, D. M. McGarrell, Y. Sun, C. T. Brown, A. Porras-Alfaro, C. R. Kuske, and J. M. Tiedje, "Ribosomal Database Project: data and tools for high throughput rRNA analysis," *Nucleic acids research*, p. gkt1244, 2013.

[36] C. Quince, A. Lanzén, T. P. Curtis, R. J. Davenport, N. Hall, I. M. Head, L. F. Read, and W. T. Sloan, "Accurate determination of microbial diversity from 454 pyrosequencing data," *Nature methods*, vol. 6, no. 9, pp. 639–641, 2009.

[37] W. Li and A. Godzik, "Cd-hit: a fast program for clustering and comparing large sets of protein or nucleotide sequences," *Bioinformatics*, vol. 22, no. 13, pp. 1658–1659, 2006.

[38] E. Stackebrandt and B. Goebel, "Taxonomic note: a place for DNA-DNA reassociation and 16S rRNA sequence analysis in the present species definition in bacteriology," *International Journal of Systematic Bacteriology*, vol. 44, no. 4, pp. 846–849, 1994.

[39] P. D. Schloss, S. L. Westcott, T. Ryabin, J. R. Hall, M. Hartmann, E. B. Hollister, R. A. Lesniewski, B. B. Oakley, D. H. Parks, C. J. Robinson, *et al.*, "Introducing mothur: open-source, platform-independent, community-supported software for describing and comparing microbial communities," *Applied and environmental microbiology*, vol. 75, no. 23, pp. 7537–7541, 2009.

[40] T. Thomas, J. Gilbert, and F. Meyer, "Metagenomics-a guide from sampling to data analysis," *Microb Inform Exp*, vol. 2, no. 3, pp. 1–12, 2012.

[41] F. M. J. Willems, "The context-tree weighting method: extensions," *IEEE Trans. Inf. Theory*, vol. IT-44, no. 2, pp. 792–798, 1998.